

11-16-00

A

PTO/SB/05 (12/97)

Please type a plus sign (+) inside this box → ☐

Approved for use through 09/30/00. OMB 0651-0032

Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number

UTILITY PATENT APPLICATION TRANSMITTAL

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Attorney Docket No	11-29694
First Named Inventor or Application Identifier	Sanjive Agarwala, et al.
Title	Write Allocation Counter for Transfer Controller With Hub and Ports
Express Mail Label No	EL547741879

APPLICATION ELEMENTS

See MPEP Chapter 600 concerning utility patent application contents

ADDRESS TO:

Assistant Commissioner for Patents
Box Patent Application
Washington, DC 20231

1. <input checked="" type="checkbox"/> *Fee Transmittal Form (e.g., PTO/SB/17) (Submit an original, and a duplicate for fee processing)	6. <input type="checkbox"/> Microfiche Computer Program (Appendix)
2. <input checked="" type="checkbox"/> Specification [Total Pages 25] (preferred arrangement set forth below) - Descriptive title of the Invention - Cross References to Related Applications - Statement Regarding Fed sponsored R&D - Reference to Microfiche Appendix - Background of the Invention - Brief Summary of the Invention - Brief Description of the Drawings (if filed) - Detailed Description - Claim(s) - Abstract of the Disclosure	7. <input type="checkbox"/> Nucleotide and/or Amino Acid Sequence Submission (if applicable, all necessary) a. <input type="checkbox"/> Computer Readable Copy b. <input type="checkbox"/> Paper Copy (identical to computer copy) c. <input type="checkbox"/> Statement verifying identical of above copies
3. <input checked="" type="checkbox"/> Drawing(s) (35 USC d113) [Total Sheets 5]	<h2>ACCOMPANYING APPLICATION PARTS</h2>
4. Oath or Declaration [Total Pages 2] a. <input checked="" type="checkbox"/> Newly Executed (original or copy) b. <input type="checkbox"/> Copy from a prior application (37 CFR §1.63(d)) (for continuation/divisional with Box 17 completed) [Note Box 5 below] i. <input type="checkbox"/> DELETION OF INVENTOR(S) Signed statement attached deleting inventor(s) named in the prior application, see 37 CFR §1 63(d)(2) and 1 33(b)	
5. <input type="checkbox"/> Incorporation By Reference (useable if Box 4b is checked) The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.	
17. If a CONTINUING APPLICATION, check appropriate box and supply the requisite information below and in a preliminary amendment: <input type="checkbox"/> Continuation <input type="checkbox"/> Divisional <input type="checkbox"/> Continuation-in-part (CIP) of prior application No: / Prior application information: Examiner _____ Group / Art Unit: _____	

*A new statement is required to be entitled to pay small entity fees, except where one has been filed in a prior application and is being relied upon

18. CORRESPONDENCE ADDRESS

☒ Customer Number or Bar Code Label

23494

or Correspondence address below

(Insert Customer No. or Attach bar code label here)

NAME	Robert D. Marshall, Jr.		
ADDRESS			
CITY	STATE	ZIP CODE	
COUNTRY	TELEPHONE	FAX	

Name (Print/Type)	Robert D. Marshall, Jr.	Registration No. (Attorney/Agent)	28,527
Signature	<i>Robert D. Marshall, Jr.</i>	Date	November 15, 2000

Burden Hour Statement. This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Box Patent Application, Washington, DC 20231.

11/15/00

09/15/00

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number

FEE TRANSMITTAL

Patent fees are subject to annual revision on October 1.

These are the fees effective October 1, 1997

Small Entity payments must be supported by a small entity statement, otherwise large entity fees must be paid. See Forms PTO/SB/09-12

Express Mailing Label No.: **EL547741879**TOTAL AMOUNT OF PAYMENT (\$ **710.00**)**Complete If Known**

Application Number

Filing Date

November 15, 2000

First Named Inventor

Sanjive Agarwala, et al.

Examiner Name

Group / Art Unit

Attorney Docket No.

TI-29694**METHOD OF PAYMENT**

- 1.
- ☒
- The Commissioner is hereby authorized to charge to the following Deposit Account,

Deposit Account Number

20-0668

Deposit Account Name

Texas Instruments Incorporated

- ☒
- Charge any additional fee required or credit any overpayment

- ☐
- Charge all indicated fees and any additional fee required or credit any overpayment

- 2.
- ☐
- Payment Enclosed:**



Check



Money Order



Other

FEE CALCULATION**1. BASIC FILING FEE**

Large Fee Code	Entity Fee (\$)	Small Fee Code	Entity Fee (\$)	Fee Description	Fee Paid
101	710	201	395	Utility filing fee	\$ 710
106	310	206	165	Design filing fee	\$
107	480	207	270	Plant filing fee	\$
108	760	208	395	Reissue filing fee	\$
114	150	214	75	Provisional filing fee	\$
SUBTOTAL (1)					(\$ 710)

2. EXTRA CLAIM FEES

Total Claims	Extra Claims	Fee from below	Fee Paid
14	-20**= 0 x 18 = 0		
Independent Claims 1	-3** = 0 x 80 = 0		
Multiple Dependent	260 =		

**or number previously paid, if greater, For Reissue, see below

Large Fee Code	Entity Fee (\$)	Small Fee Code	Entity Fee (\$)	Fee Description
103	18	203	11	Claims in excess of 20
102	78	202	41	Independent Claims in excess of 3
104	260	204	135	Multiple dependent claims in excess of 3
109	78	209	41	**Reissue independent claims over original patent
110	18	210	11	**Reissue claims in excess of 20 and over original patent
SUBTOTAL (2)				(\$ 0

FEE CALCULATION (continued)**3. ADDITIONAL FEES**

Large Fee Code	Entity Fee (\$)	Small Fee Code	Entity Fee (\$)	Fee Description	Fee Paid
105	130	205	65	Surcharge - late filing fee	
127	50	227	25	Surcharge - late provisional filing fee or cover sheet	
139	130	139	130	Non-English specification	
147	2,520	147	2,520	For filing a request for reexamination	
112	920*	112	920*	Requesting publication of SIR prior to Examiner action	
113	1,840*	113	1,840*	Requesting publication of SIR after Examiner action	
115	110	215	55	Extension for reply within first month	
116	380	216	200	Extension of time within second month	
117	870	217	475	Extension of time within third month	
118	1,360	218	755	Extension of time within fourth month	
128	1,850	228	1,030	Extension of time within fifth month	
119	300	219	155	Notice of Appeal	
120	300	220	155	Filing a brief in support of an appeal	
121	260	221	135	Request for oral hearing	
138	1,510	138	1,510	Petition to institute a public use proceeding	
140	110	240	55	Petition to revive - unavoidable	
141	1,210	241	660	Petition to revive - unintentional	
142	1,210	242	660	Utility issue fee (or reissue)	
143	430	243	225	Design issue fee	
144	580	244	335	Plant issue fee	
122	130	122	130	Petitions to the Commissioner	
123	50	123	50	Petitions related to provisional applications	
126	240	126	240	Submission of Information Disclosure Stmt	
581	40	581	40	Recording each patent assignment per properly (time number of properties)	
146	760	246	395	Filing a submission after final rejection (37 CFR 1.129(a))	
149	760	249	395	For each additional invention to be examined (37 CFR 1.129(b))	

Other fee (specify)

Other fee (specify)

*Reduced by Basic Filing Fee Paid

SUBTOTAL (3)

SUBMITTED BY

Typed or Printed Name

Robert D. Marshall, Jr.

Signature

Robert D. Marshall, Jr.

Date

November 15, 2000

Complete (if applicable)

Reg Number

28,527

Deposit Account User ID

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re the Application of:

TI-29694

Sanjive Agarwala, et al.

Serial No:

Filed: November 15, 2000

For: Write Allocation Counter for Transfer Controller With Hub and Ports

PRELIMINARY AMENDMENT

Ass't Commissioner for Patents
Washington, DC 20231

Dear Sir:

EXPRESS MAILING" Mailing Label No. EL547741879. Date of Deposit: November 15, 2000. I hereby certify that this paper is being deposited with the U.S. Postal Service Express Mail Post Office to Addressee Service under 37 CFR 1.10 on the date shown above and is addressed to: Ass't Commissioner for Patents, Washington, D.C. 20231.


Robin E. Barnum

Please amend the specification by inserting before the first line, the following sentence:

--This application claims priority under 35 USC §119(e)(1) of Provisional Application Number 60/169,415, filed December 7, 1999.--

Respectfully submitted,



Robert D. Marshall, Jr.
Attorney for Applicants
Reg. No. 28,527

Texas Instruments Incorporated
P.O. Box 655474, MS 3999
Dallas, TX 75265
(972) 917-5290

0051343-1500
ELECTRONIC MAILING

WRITE ALLOCATION COUNTER FOR
TRANSFER CONTROLLER WITH HUB AND PORTS

Sanjive Agarwala
Iain Robertson
David A. Comisky
Charles L. Fuoco

TECHNICAL FIELD OF THE INVENTION

The technical field of this invention is digital signal processing and more particularly control of data transfers within a digital signal processing system.

5

BACKGROUND OF THE INVENTION

Digital signal processing (DSP) differs significantly from general purpose processing performed by micro-controllers and microprocessors. One key difference is the strict requirement for real time data processing. For example, in a modem application, it is absolutely required that every sample be processed. Even losing a single data point might cause a digital signal processor application to fail. While processing data samples may still take on the model of tasking and block processing common to general purpose processing, the actual data movement within a digital signal processor system must adhere to the strict real-time requirements of the system.

10

15

20

25

As a consequence, digital signal processor systems are highly reliant on an integrated and efficient direct memory access (DMA) engine. The direct memory access controller is responsible for processing transfer requests from peripherals and the digital signal processor itself in real time. All data movement by the direct memory access must be capable of occurring without central processing unit (CPU) intervention in order to meet the real time requirements of the system. That is, because the CPU may operate in a software tasking model where scheduling of a task is not as tightly controlled as the data streams the tasks operate on require, the direct

memory access engine must sustain the burden of meeting all real time data stream requirements in the system.

The early direct memory access has evolved into several successive versions of centralized transfer controllers and more recently into the transfer controller with hub and ports architecture. The transfer controller with hub and ports architecture is described in U.K. Patent Application No. 9909196.9 filed April 10, 1999 entitled "TRANSFER CONTROLLER WITH HUB AND PORTS ARCHITECTURE" (TI-28983).

A first transfer controller module was developed for the TMS330C80 digital signal processor from Texas Instruments. The transfer controller consolidated the direct memory access function of a conventional controller along with the address generation logic required for servicing cache and long distance data transfer, also called direct external access, from four digital signal processors and a single RISC (reduced instruction set computer) processor.

The transfer controller architecture of the TMS330C80 is fundamentally different from a direct memory access in that only a single set of address generation and parameter registers is required. Prior direct memory access units required multiple sets for multiple channels. The single set of registers, however, can be utilized by all direct memory access requestors. Direct memory access requests are posted to the transfer controller via set of encoded inputs at the periphery of the device. Additionally, each of the digital signal processors can submit requests to the transfer controller. The external encoded inputs are called "externally initiated packet transfers" (XPTs). The digital signal processor initiated transfers are referred to as "packet

transfers" (PTs). The RISC processor could also submit packet transfer requests to the transfer controller.

The transfer controller with hub and ports introduced several new ideas concepts. The first was uniform pipelining.

5 New digital signal processor devices containing a transfer controller with hub and ports architecture have multiple external ports, all of which look identical to the hub. Thus peripherals and memory may be freely interchanged without affecting the hub. The second new idea is the concept of
10 concurrent execution of transfers. That is, up to N transfers may occur in parallel on the multiple ports of the device, where N is the number of channels in the transfer controller with hub and ports core. Each channel in the transfer controller with hub and ports core is functionally just a set
15 of registers. This set of registers tracks the current source and destination addresses, the word counts and other parameters for the transfer. Each channel is identical, and thus the number of channels supported by the transfer controller with hub and ports is highly scalable.

20 Finally the transfer controller with hub and ports includes a mechanism for queuing transfers up in a dedicated queue memory. The TMS320C80 transfer controller permitted only was one transfer outstanding per processor at a time. Through the queue memory provided by the transfer controller
25 with hub and ports, processors may issue numerous transfer requests up to the queue memory size before stalling the digital signal processor.

SUMMARY OF THE INVENTION

5 The transfer controller with hub and ports has undergone significant refinements in implementation that followed the original description in U.K. Patent Application No. 9909196.9 filed April 10, 1999 entitled "TRANSFER CONTROLLER WITH HUB AND PORTS ARCHITECTURE." One such refinement is the use of a write allocation counter in the source pipeline. The algorithm upon which this write allocation counter operates is also key to the invention.

10 The write allocation counter is used to throttle the fast source port read operations based on the amount of data that can be consumed immediately by the slow destination port write reservation station and the channel data router buffers. This ensures that source port response queue is not blocked with
15 data that cannot be consumed by the requesting channel data router and its slow destination port and thereby blocking out source port from providing data to the other channel destination ports.

BRIEF DESCRIPTION OF THE DRAWINGS

These and other aspects of this invention are illustrated in the drawings, in which:

5 Figure 1 illustrates in a functional block diagram the basic principal features of the transfer controller with hub and ports architecture and related functions;

Figure 2 illustrates the queue manager interface to the transfer controller hub unit;

10 Figure 3 illustrates the transfer controller source and destination operational pipelines;

Figure 4 illustrates the source pipeline P-stage including the write allocation counter and write allocation counter rules unit; and

15 Figure 5 illustrates the write allocation counter algorithm in flow chart form.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

5 The transfer controller with hub and ports transfer controller with hub and ports architecture is optimized for efficient passage of data throughout a digital signal processor chip. Figure 1 illustrates a block diagram of the principal features of the transfer controller with hub and ports. It consists of a system of a single hub 100 and multiple ports 111 through 115.

10 The transfer controller with hub and ports functions in conjunction with a transfer request bus having a set of nodes 117, which bring in transfer request packets at input 103. These transfer request bus nodes individually receive transfer requests packets from transfer requestors 116 which are processor-memory nodes or other on-chip functions which send and receive data.

15 Secondly, the transfer controller uses an additional bus, the data transfer bus having a set of nodes 118, to read or write the actual data at the requestor nodes 116. The data transfer bus carries commands, write data and read data from a special internal memory port 115 and returns read data to the transfer controller hub via the data router 150 at inputs 104.

20 The transfer controller has, at its front-end portion, a request queue manager 101 receiving transfer requests in the form of transfer request packets at its input 103. Request queue manager 101 prioritizes, stores and dispatches these as required.

25 Request queue manager 101 connects within the transfer controller hub unit 100 to the channel request registers 120 which receive the data transfer request packets and process

them. In this process, request queue manager 101 first prioritizes the transfer request packets and assigns them to one of the N channel request registers 120. Each of the N channel request registers 120 represents a priority level.

5 If there is no channel available for direct processing of the transfer request packet, it is stored in the queue manager memory 102. Queue manager memory 102 is preferably a random access memory (RAM). The transfer request packet is then assigned at a later time when a channel becomes available.

10 The channel registers interface with the source 130 and destination 140 control pipelines which effectively are address calculation units for source (read) and destination (write) operations.

15 Outputs from these pipelines are broadcast to M ports through the transfer controller ports I/O subsystem 110. I/O subsystem 110 includes a set of hub interface units, which drive the M possible external ports units. Four such external ports are shown in Figure 1 as external ports 111 through 114. The external ports units (also referred to as application
20 units) are clocked either at the main processor clock frequency or at a different external device clock frequency. The external device clock frequency may be lower than or higher than the main processor clock frequency. If a port operates at its own frequency, synchronization to the core
25 clock is required.

30 As an example of read-write operations at the ports, consider a read from external port node 112 followed by a write to external port node 114. First the source pipeline addresses port 112 for a read. The data is returned to the transfer controller hub through the data router 150. On a

later cycle the destination control pipeline addresses port 114 and writes the data at port 114. External ports as described here do not initiate transfer requests but merely participate in reads and writes requested elsewhere on the chip. Read and write operations involving the processor-memory (transfer requestors) nodes 116 are initiated as transfer request packets on the transfer request bus 117. The queue manager 101 processes these as described above. On a later cycle a source pipeline output (read command/address) is generated which is passed at the internal memory port to the data transfer bus 118 in the form of a read. This command proceeds from one node to the next in pipeline fashion on the data transfer bus. When the processor node addressed is reached, the read request causes the processor-memory node to place the read data on the bus for return to the data router 150. On a later cycle, a destination pipeline output passes the corresponding write command and data to the internal memory port and on to the data transfer bus for writing at the addressed processor node.

The channel parameter registers 105 and port parameters registers 106 hold all the necessary parametric data as well as status information for the transfer controller hub pipelines to process the given transfer. Both pipelines share some of the stored information. Other portions relate specifically to one pipeline or the other.

Figure 2 illustrates the interface of request queue manager 101 to the transfer controller hub unit boundary and particularly the request queue manager communications with the channel request registers 200, channel parameter registers 105 and port parameters registers 106. Channel parameters

registers 105 and port parameters registers 106 store critical data regarding for example, types of transfers, mode information, status, and much other information critical to the transfer process.

Channel request registers 200 pass information used in the source control pipeline 130 for generation of the read/pre-write commands 221. Similarly, channel request registers 200 pass information used in the destination control pipeline 140 for the generation of write command/write data words 222. Read response data 104 from the ports is returned to the destination pipeline via the data router 150.

Figure 3 illustrates the possible pipelines in a transfer controller implementation. Table 1 shows the particular tasks performed during the pipeline stages in the preferred embodiment. In specific implementations, one or more stages may be combined but the tasks for the individual pipeline stages are essentially as shown in Table 1.

Pipeline Stage	Function
Q	Interrogates state of queues within ports
M	Maps port ready signals to channels
P	Prioritize highest priority channel with ready ports
A0	First half of address update cycle
A1	Second half of address update cycle
C	Issues command to ports

Table 1

The channel request registers 200 pass information used in the source pipeline stages 301 to 306 for generation of the read/pre-write commands 221. Similarly, the channel request registers 200 pass information used in the destination pipeline stages 311 to 315 for the generation of write command/write data words 222. Read response data 104 from the ports is returned to the destination pipeline via the data router 150.

SOURCE PIPELINE

The functions of the first three individual stages of the source pipeline (Q, M and P stages) may be combined into just two stages, the M-stage and the P-stage. The P-stage is of the most significance in the present invention. The source pipeline functions perform all the source reads and destination reservation station pre-writes for a write driven processing transfer. Write driven processing is performed when the read port is relatively slow as compared to the destination port. The pre-allocated write space in the ports is divided into multiples of write burst size and each such entry is termed as a reservation station. A pre-write reserves a reservation station entry in a port for the channel initiating the request.

P-stage 303 illustrated in Figure 4 starts with two parallel computations. On a per channel basis, ready feedback unit 401 computes source and destination port availability based on queue counters, indicated by read queue counter 424 and write queue counter 425, and previous clock port usage in the P-stage. In the second computation, write allocation counter/rules unit 402 applies write allocation counter rules

and determines if read or pre-write needs to be done per channel basis. As noted below, the write allocation counter algorithm is dependent upon the source default read burst size, which is indicated by source default size signal 427.

5 Both of these results are fed to the channel ready evaluation unit 403. Channel ready evaluation unit 403 maps port availability to the type of operation (read/pre-write) and determines if the channel can be scheduled for processing. Channel ready evaluation unit 403 also receives a read/not
10 write driven signal 426 indicating whether read driven or write driven processing is used.

The output the channel ready evaluation unit 403 indicates all channels which are ready for command processing. Next the prioritizer unit 406 selects the highest priority
15 channel which will be scheduled for processing in the following stages. The selected channel is used as control for a multiplexer to determine the port number on which the command will be scheduled.

Source pipeline P-stage 303 prioritizes and schedules one
20 of the active channels for command processing in the following A0/A1 stages of pipeline. The output of source pipeline P-stage 303 is the selected channel number 411 and the port number 412 to which the command will be scheduled. A valid read/write command is identified by either a channel read
25 valid signal 408 or a write valid signal 410 being asserted. A valid pre-write command is identified by a pre-write valid signal 409.

There are four macro-level logical operations involved in source pipeline P-stage 303 operation. The ready feedback
30 unit 401 further qualifies the channel SRC ready signal 422

and DST ready signal 423 from source pipeline M-stage 302 to alleviate one cycle inaccuracy in the local port counter values due to the pipelining effect. Write allocation counter/rules unit 402 identifies if a read or reservation station pre-write operation needs to be performed for channels performing write-driven processing. Channel ready evaluation unit 403 determines if a channel is ready for scheduling based on the source/destination port availability and read/pre-write type of operation. Prioritizer unit 406 selects the highest priority channel that is ready for scheduling.

WRITE ALLOCATION COUNTER (WAC)

The source pipeline has a write allocation counter/rules unit 402 associated with each and every channel. Each write allocation counter/rules unit 402 sequences write driven processing operations. This counter is used to throttle fast source port read operations based on the amount of data that can be consumed immediately by a slow destination port write reservation station and the channel data router 150 buffers. This ensures that a source port response queue is not blocked with data that cannot be consumed by the requesting channel data router 150 and its corresponding slow destination port. This condition blocks the source port from providing data to the other channel destination ports.

A reservation station pre-write at the destination port causes write allocation counter/rules unit 402 to increment the write allocation count by the smallest of the burst size of the destination port, the number of words left to transfer or the size based on alignment. A non-zero value of the write allocation count enables primary reads at the source port.

The write allocation count is decremented as reads get posted to the source port provided the write allocation count does not go below zero. Note that it is possible to do multiple reservation station writes to the destination port before performing reads at the source port.

WRITE ALLOCATION COUNTER ALGORITHM

For a write driven process, the decision to issue a read to the source port or a reservation station pre-write to the destination port is made based on the value of the write allocation count. For maximum performance, it would be desirable to do read/write operations of burst size and avoid read/write operations of smaller than burst size. The following rules reflect this reasoning and bias toward reservation station writes until the write allocation count is greater than read burst size or the reservation station is full. The rules are applied by the write allocation counter/rules unit 402 which implements the write allocation counter algorithm.

To clarify the write allocation counter algorithm, consider the concept from its basic requirements. In write driven processing the transfer controller hub is not allowed to read from a source port more than the amount of data for which space has been allocated in the destination port. A counter is needed to keep track of how much space has been pre-allocated. This is the write allocation count.

Figure 5 illustrates the write allocation counter algorithm. The write allocation count is incremented by the number of words allocated in the reservation station by a pre-write (511). The write allocation counter is decremented by

the number of words read during a read (505). The value of the write allocation count at any time is the number of words which can be read without causing a backlog in the source port. Thus the write allocation count is initialized to the capacity of the data router 150 R (512).

5 A pre-write cannot occur unless there is at least one reservation station entry available in the destination port. Thus a pre-write may have to wait for a destination write from a reservation station to free space. A pre-write accomplishes two things. A pre-write increments (508) the write allocation count by the number of words being reserved in the destination port's reservation station. A pre-write causes the hub's local counter of number of available reservation station entries for the destination port to decrement. Thus the count of available reservation station entries decrements. During the source pipeline C-Stage 306 operation at the end of a pre-write, the address alignment, word count, and channel number of the pre-write are sent to the destination port. These quantities are stored in the reservation station entry assigned to the destination port. Sometime later, the data which has been read is transferred to the destination port. When all the data for a particular reservation station entry has arrived, the destination port can complete the write. It then marks that reservation station entry as not in use, and sends a signal to the hub telling it to increment its count of available reservation station entries associated with that port.

25 A separate count of available reservation station entries for each port is maintained in the ready feedback unit 401 of Figure 4. Ready feedback unit 401 contains one of such

counter for each port. When the counter for a specific port decrements to zero, the condition "reservation station full" becomes true for that port.

Because some amount of data can be staged in the data router 150, it is possible to read slightly more than the amount pre-allocated before stalling the source port. If the data router 150 holds R words, it is possible to read R words more than has been pre-allocated in destination port reservation stations without causing a backlog in the source port. It is desirable to allow this to occur to ensure reads of up to the burst size can be achieved wherever possible. If the initial value of the write allocation count is zero, this would require support for negative numbers. To avoid this and to simplify the implementation, in the preferred embodiment the write allocation count is pre-initialized (512) to R, the number of words which can be staged in the data router 150. Using this technique, the minimum value the write allocation count can take is 0. Any read which would cause the write allocation count to fall below 0 will be annulled.

A read should be performed only if the write allocation counter is greater than or equal to the read burst size. If the read burst size is larger than the write burst size, it will take several pre-writes before the write allocation count is large enough for a read to start.

The decision about what process proceeds next is carried out in the source pipeline P-stage 303. However, the write allocation count value is maintained in A-stages. Therefore, at the point at which a decision is made based on the value of write allocation counter, there may be changes to the value pending as a result of requests scheduled on previous cycles.

These need to be accounted for in order to reduce the probability of having to annul a read. If there was no write in the previous cycle, then the value of write allocation count is not going to be increased in the coming clock cycle.

5 It may be about to be decreased in the coming clock cycle by a previously scheduled read or may not be decreased if no read is scheduled. In this case, it is desirable to minimize the probability of a scheduled read having to be annulled due to an attempt to decrement the write allocation counter below
10 zero by further requiring that the write allocation count be greater than or equal to the data storage capacity of data router 150.

Figure 5 illustrates the write allocation counter algorithm in flow chart form. The write allocation counter
15 (510) is initialized at the capacity R of data router 150 (512). The write allocation counter algorithm operates as follows:

Do a source data read (block 505) if:

- 20 (1) the write allocation count is greater than or equal to the read burst size (Yes at decision block 502)
AND
there was pre-write on previous cycle (Yes at decision block 504);

OR

- 25 (2) the write allocation count is greater than or equal to the read burst size (Yes at decision block 502)
AND
there was not pre-write on previous cycle (No at decision block 504), AND
30 the write allocation count is greater than the

capacity R of data router 150 (Yes at decision block 501);

OR

- 5 (3) the write allocation count is not greater than or equal to the read burst size (No at decision block 502) AND
the reservation station full (Yes at decision block 503).

10 Do a pre-write (block 511) if:

- (1) the write allocation count is not greater than or equal to the read burst size (No at decision block 502), AND
the reservation station is not full (No at decision block 503);

OR

- (2) the write allocation count is greater than or equal to the read burst size (Yes at decision block 502), AND
there was not a pre-write on previous cycle (No at decision block 504), AND
the write allocation counter is not greater than or equal to the data capacity R of data router 150 (No at decision block 501), AND
the reservation station is not full (No at decision block 503).

The write allocation count is decremented on a source data read (block 508). The write allocation count is incremented on a pre-write (block 509). These rules were

selected empirically through simulation. They provide a good compromise between low latency and efficiency. High latency would result if the rule were to wait for write allocation counter to fully resolve and allow new requests to start only every few cycles. Lack of efficiency would result if too many reads were annulled because the write allocation counter would go below zero.

A additional refinement adjusts the data transfer size. If a read is annulled because the write allocation counter would otherwise have gone negative and the write reservation station is full, then the value of write allocation counter is used as the read transfer size when the read is re-attempted. If this is not done then deadlock would result. Reads could continue to be attempted and annulled until some space was freed up in the reservation station. By the same token, however, no space would actually be freed up because the data for the pre-write would never arrive because the read would never actually take place. The adjustment in the data transfer size prevents this deadlock.

WHAT IS CLAIMED IS:

097349-11500

1 1. A method for tracking allocated space in a write
2 reservation station of a data transfer controller using a
3 write allocation count, said method comprising the steps of:
4 initializing said write allocation count prior to
5 performance of any data transfers;
6 incrementing said write allocation count on allocation of
7 a block of write reservation station space at a data
8 destination;
9 decrementing said write allocation count on a read from
10 a data source;
11 if said write allocation count meets predetermined
12 criteria, then reading from said data source, transferring
13 said read data to a data destination via a data routing
14 channel and storing said transferred data in allocated
15 reservation station space; and
16 if said write allocation count does not meet said
17 predetermined criteria, then performing no further allocations
18 of space to said write reservations station until said write
19 allocation count meets said predetermined criteria.

1 2. The method of claim 1, wherein:
2 said predetermined constant of said step of initializing
3 said write allocation count equals a number of data words
4 storable in said data routing channel.

1 3. The method of claim 1, wherein:
2 said step of incrementing said write allocation count on
3 allocation of a block of write reservation station space

4 increments said write allocation counter by an amount equal to
5 a number of data words allocated.

1 4. The method of claim 1, wherein:
2 said step of decrementing said write allocation count on
3 a read from a data source decrements said write allocation
4 counter by an amount equal to a number of data words read.

1 5. The method of claim 1, wherein:
2 said step of reading from said data source reads data in
3 an amount equal to a read burst size constant related to a
4 default read burst size of said data source.

1 6. The method of claim 5, wherein:
2 said predetermined criteria of said write allocation
3 count includes whether said write allocation count is
4 greater than or equal to said read burst size constant.

1 7. The method of claim 5, wherein:
2 said predetermined criteria of said write allocation
3 count includes whether said write allocation count is greater
4 than or equal to a number of data words storable in said data
5 routing channel.

1 8. The method of claim 5, wherein:
2 said predetermined criteria of said write allocation
3 count is met if
4 said write allocation count is greater than or
5 equal to said read burst size constant, and

6 an allocation of a block of write reservation
7 station space was made in an immediately prior cycle.

1 9. The method of claim 5, wherein:

2 said predetermined criteria of said write allocation
3 count is met if

4 said write allocation count is greater than or
5 equal to said read burst size constant, and

6 an allocation of a block of write reservation
7 station space was not made in an immediately prior cycle,
8 and

9 said write allocation count is greater than or
10 equal to a number of data words storable in said data
11 routing channel.

1 10. The method of claim 5, wherein:

2 said predetermined criteria of said write allocation
3 count is met if

4 said write allocation count is not greater than or
5 equal to said read burst size constant, and

6 all write reservation station space at said data
7 destination has been allocated.

1 11. The method of claim 5, where:

2 said predetermined criteria of said write allocation
3 count is not met if

4 said write allocation count is not greater than or
5 equal to said read burst size constant, and

6 all write reservation station space at said data
7 destination have not been allocated.

1 12. The method of claim 5, wherein:
2 said predetermined criteria of said write allocation
3 count is not met if
4 said write allocation count is greater than or
5 equal to said read burst size constant, and
6 an allocation of a block of write reservation
7 station space was not made in an immediately prior cycle,
8 and
9 said write allocation count is not greater than or
10 equal to a number of data words storable in said data
11 routing channel, and
12 all write reservation station space at said data
13 destination have not been allocated.


1 13. The method of claim 1, further comprising the steps
2 of:
3 reading data from said reservation station space and
4 writing said read data to said data destination at rate
5 determined by said data destination;
6 deallocating a block of write reservation space at said
7 data destination upon reading data from said reservation
8 station space and writing said read data to said data
9 destination; and
10 said step of incrementing said write allocation count on
11 allocation of a block of write reservation station space at
12 said data destination occurs only if at least some write
13 reservation station space has not been allocated.

1 14. The method of claim 13, wherein:
2 said step of reading data from said reservation station
3 space reads data in an amount equal to a write burst size
4 constant related to a default write burst size of said data
5 destination; and
6 said step of deallocating a block of write reservation
7 space at said data destination deallocates a block having a
8 size equal to said write burst size constant.

ABSTRACT

1 The transfer controller with hub and ports uses a write
2 allocation counter and algorithm to control data reads from a
3 source port. The write allocation count is the amount of data
4 that can be consumed immediately by the write reservation
5 station of a slow destination port and the channel data router
6 buffers. This is used to throttle fast source port read
7 operations to whole read bursts until space to adsorb the read
8 data is available. This ensures that the source port response
9 queue is not blocked with data that cannot be consumed by the
10 channel data router and the slow destination port. This
11 condition would otherwise block a fast source port from
12 providing data to the other destination ports.

FIG. 1






FIG. 2

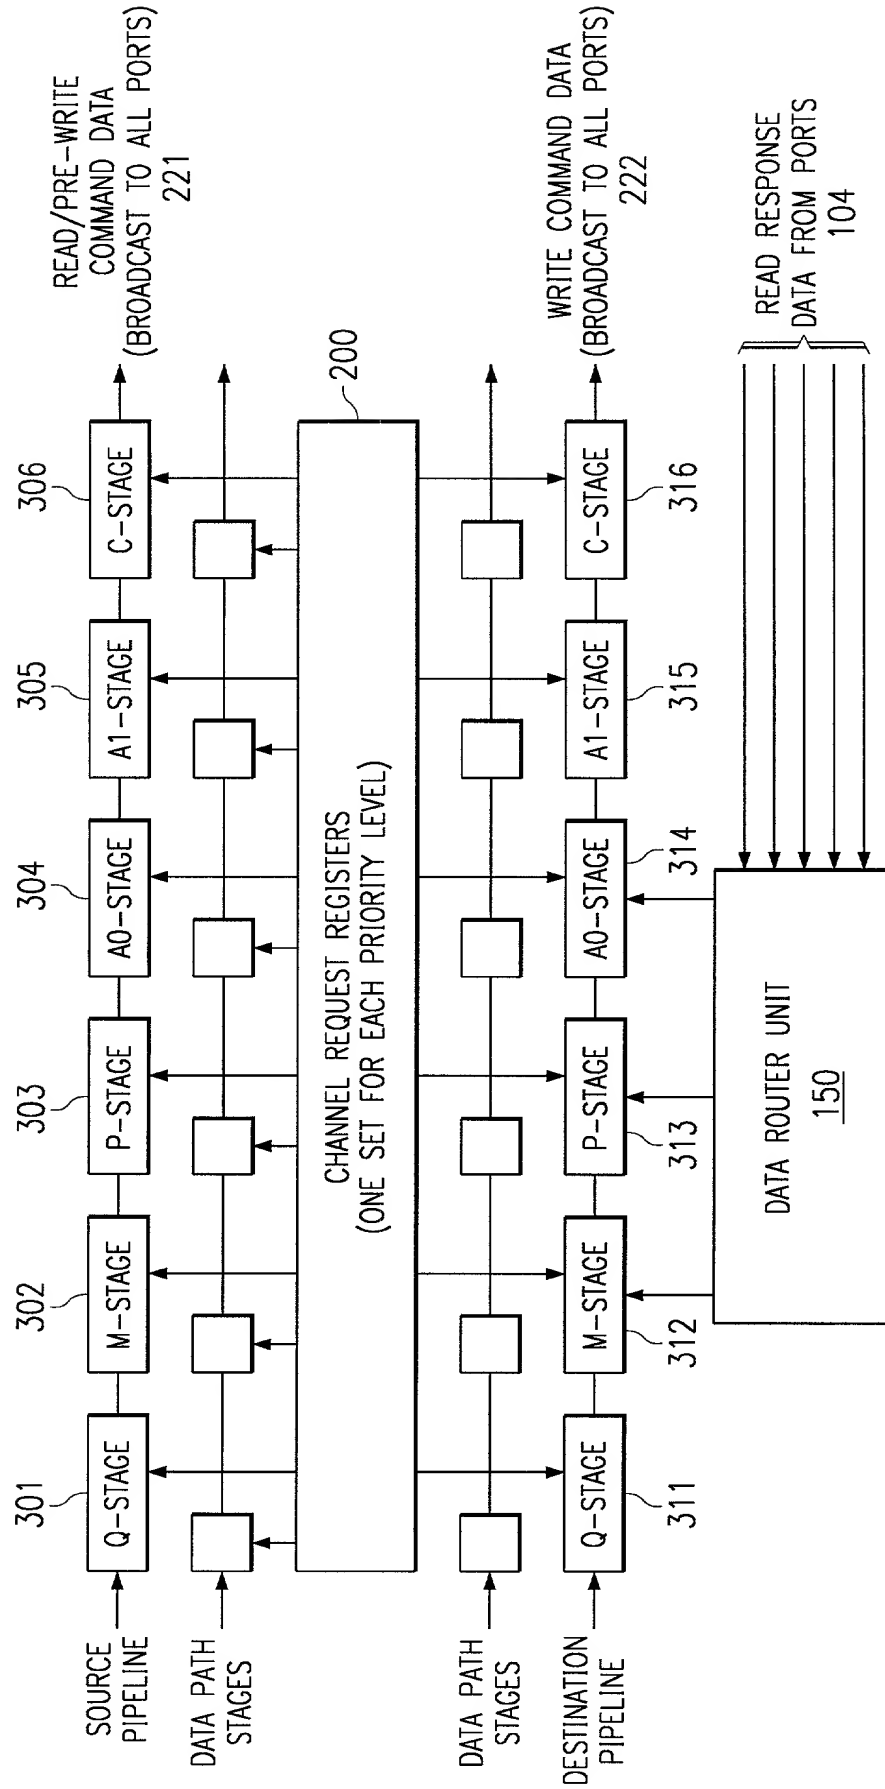


FIG. 3

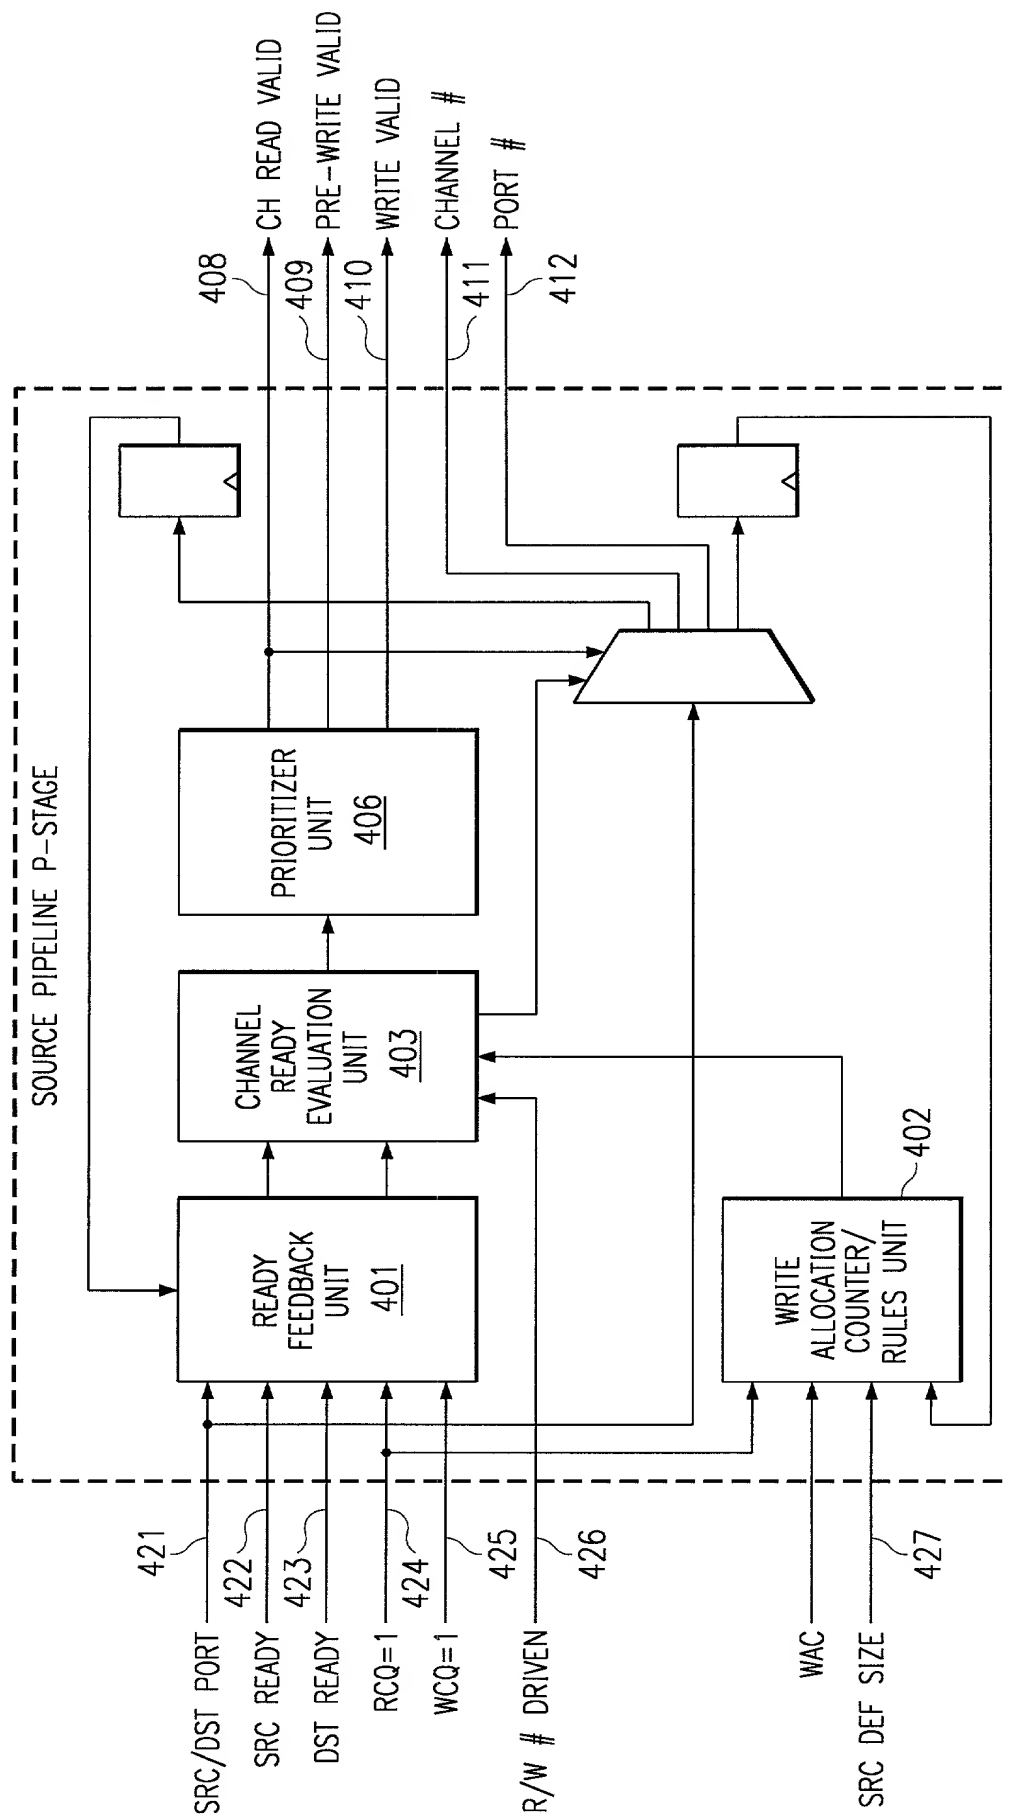


FIG. 4

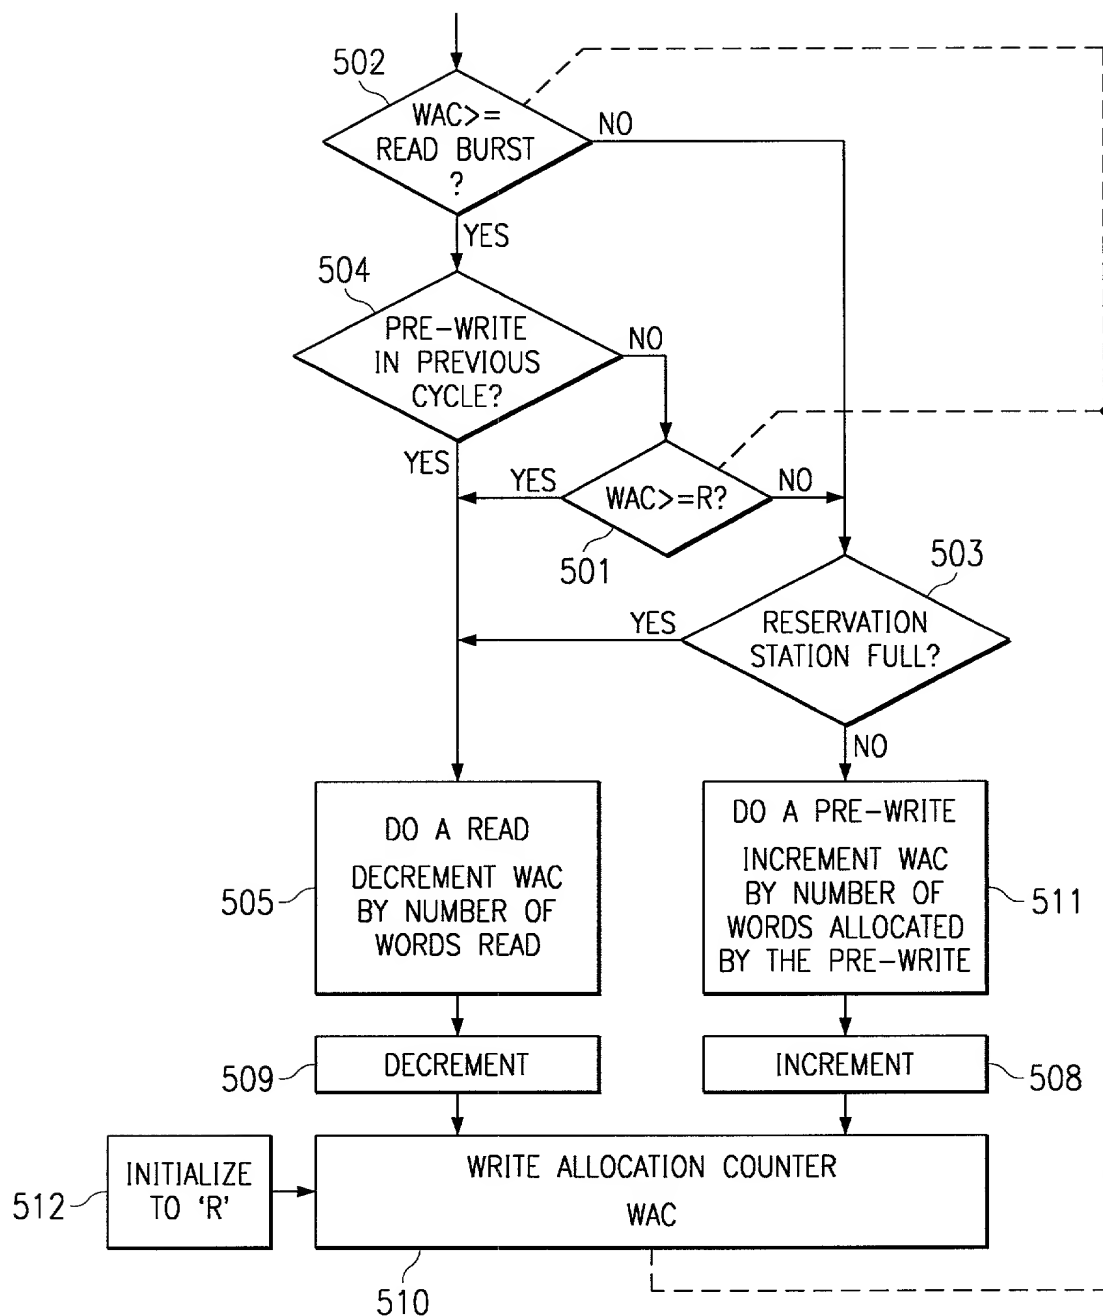


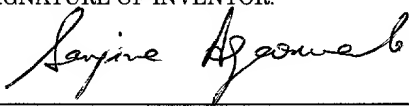
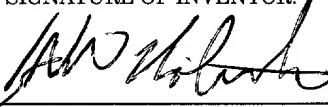

FIG. 5

PAGE 1 OF 2

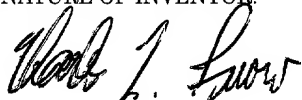

**APPLICATION FOR UNITED STATES PATENT
DECLARATION AND POWER OF ATTORNEY**

As a below named inventor, I declare that my residence, post office address and citizenship are as stated below next to my name; that I verily believe that I am the original, first and sole inventor if only one name is listed below, or an original, first and joint inventor if plural inventors are named below, of the subject matter which is claimed and for which a patent is sought on the invention entitled as set forth below, which is described in the attached specification of **Application Serial No. 60/169,415, filed 12/07/99**; that I have reviewed and understand the contents of the specification, including the claims, as amended by any amendment specifically referred to in the oath or declaration; that no application for patent or inventor's certificate on this invention has been filed by me or my legal representatives or assigns in any country foreign to the United States of America; and that I acknowledge my duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, Section 1.56;

I further declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true, and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issuing thereon.

TITLE OF INVENTION: Write Allocation Counter for Transfer Controller with Hub and Ports		
POWER OF ATTORNEY: I HEREBY APPOINT THE FOLLOWING ATTORNEYS TO PROSECUTE THIS APPLICATION AND TRANSACT ALL BUSINESS IN THE PATENT AND TRADEMARK OFFICE CONNECTED THEREWITH Frederic J. Telecky, Jr., #29,979; Jay M. Cantor, #19,906; C. Alan McClure, #31,041; William B. Kempler, #28,228; Robert D. Marshall, #28,527; Gerald E. Laws, 39,268; James F. Hollander, #27,802; Robby Holland, #33,304; Rebecca Mapstone-Lake, #36,208.		
SEND CORRESPONDENCE TO: Robert D. Marshall, Jr. Texas Instruments Incorporated P.O. Box 655474, MS 3999 Dallas, TX 75265		DIRECT TELEPHONE CALLS TO: Robert D. Marshall, Jr. (972) 917-5290
NAME OF INVENTOR: (1) Sanjive Agarwala	NAME OF INVENTOR: (2) Iain Robertson	NAME OF INVENTOR: (3) David A. Comisky
RESIDENCE & POST OFFICE ADDRESS: 2608 Foxcreek Drive Richardson, Texas 75082	RESIDENCE & POST OFFICE ADDRESS: 10, Grange Lane Bedfordshire MK44 3TT	RESIDENCE & POST OFFICE ADDRESS: 812 Wynnpage Lane Plano, Texas 75075
COUNTRY OF CITIZENSHIP: India	COUNTRY OF CITIZENSHIP: United Kingdom	COUNTRY OF CITIZENSHIP: United States
SIGNATURE OF INVENTOR: x 	SIGNATURE OF INVENTOR: 	SIGNATURE OF INVENTOR: x 
DATE: x 02/22/00	DATE: 22 FEB 2000	DATE: x 02/22/00

PAGE 2 OF 2

NAME OF INVENTOR: Charles L. Fuoco	NAME OF INVENTOR: N/A	NAME OF INVENTOR: N/A
RESIDENCE & POST OFFICE ADDRESS: 401 Misty Meadow Drive Allen, Texas 75013	RESIDENCE & POST OFFICE ADDRESS:	RESIDENCE & POST OFFICE ADDRESS:
COUNTRY OF CITIZENSHIP: United States	COUNTRY OF CITIZENSHIP:	COUNTRY OF CITIZENSHIP:
SIGNATURE OF INVENTOR: x 	SIGNATURE OF INVENTOR:	SIGNATURE OF INVENTOR:
DATE: x 	DATE:	DATE:

[illegible]